VLDB 2007
33rd International Conference on Very Large Data Bases
September 23-27 2007, University of Vienna, Austria

Why You Should Run TPC-DS: A Workload Analysis

Meikel Poess
Oracle USA

Raghunath Othayoth Nambiar
Hewlett-Packard Company

Dave Walrath
Sybase Inc
(in absentia)

Agenda

- Transaction Processing Performance Council (TPC)
- Scope of TPC-DS benchmark
- TPC-DS Design Considerations
- TPC-DS Workload Analysis
- TPC-DS Metric Analysis
- Q&A

Transaction Processing Performance Council

TPC Transaction Processing Performance Council

- The TPC defines transaction processing and database benchmarks and delivers trusted results to the industry.
  - Most credible, system-level benchmark evaluation test for the server industry
  - Fulfilling the role of a “Consumer Reports” for the computing industry
  - Scores are the most requested server benchmarks in server RFPs
- Active benchmarks
  - TPC-C: Online transaction processing
  - TPC-H: Data Warehouse for ad hoc queries
  - TPC-App: Application server and web services
  - TPC-E: Online transaction processing (new)
- Benchmarks under development
  - TPC-DS: Decision Support

TPC Membership

TPC Transaction Processing Performance Council

Full Members

Associate Members

What makes the TPC unique

TPC Transaction Processing Performance Council

- TPC is the only benchmark organization that requires price-performance scores across all of its benchmarks
- All tests require full documentation of the components and applications under test, so that the test can be replicated
- The TPC requires an independent audit of results prior to publication
- Extensive oversight via fair use policies
- TPC tests the whole system performance, not just a piece
- TPC is database agnostic: Oracle, IBM DB2, Sybase, Microsoft SQL Server, NonStop SQL/MX and other databases
- TPC provides cross-platform performance comparisons, a view of processor versus real performance, technology comparisons and actual cost of performance comparisons

Objectives of TPC Benchmarks

- System and database vendors
  - Competitive analysis
  - Release to release progress
  - Technology development
- Customers
  - Cross vendor/architecture performance comparison
  - Cross vendor/architecture TCO comparison
  - Evaluate new technologies
  - Eliminate investment in in-house characterization
- Research community
  - A standard yet customizable workload

TPC's DW/DSS Benchmark History

TPC Transaction Processing Performance Council

- TPC-D - Data Warehouse (1995-1999)
- TPC-R - Data Warehouse for reporting queries (99-04)
- TPC-H - Data Warehouse for ad hoc queries (99-current)
- TPC-DS - Decision Support (target 2008)
  - Latest status and specification
    - http://www.tpc.org/tpcds/default.asp
  - Series of Presentations
    - TPC-DS, Taking Decision Support Benchmarking to the Next Level, SIGMOD 2002
    - The Making of TPC-DS, VLDB 2006
    - Why You Should Run TPC-DS: Workload Analysis, VLDB 2007

Scope of TPC DS Benchmark

- Measures generally applicable aspects of a Decision Support System
  - Examine large volume of data
  - Give answers to real-world business questions
  - Execute queries of various operational requirements
  - Generate intense activity against the database server component of a system (IO, memory, CPU, Interconnect)
  - Remain closely synchronized with source OLTP database through a periodic database maintenance function
- Provides the industry
  - An objective means of comparing
    - the performance of decision support systems
    - the price-performance of decision support systems
  - A standard yet customizable workload

Gartner Inc. showed business intelligence (BI) as a top priority for CIOs.

http://www.gartner.com/2_events/conferences/bie7i.jsp

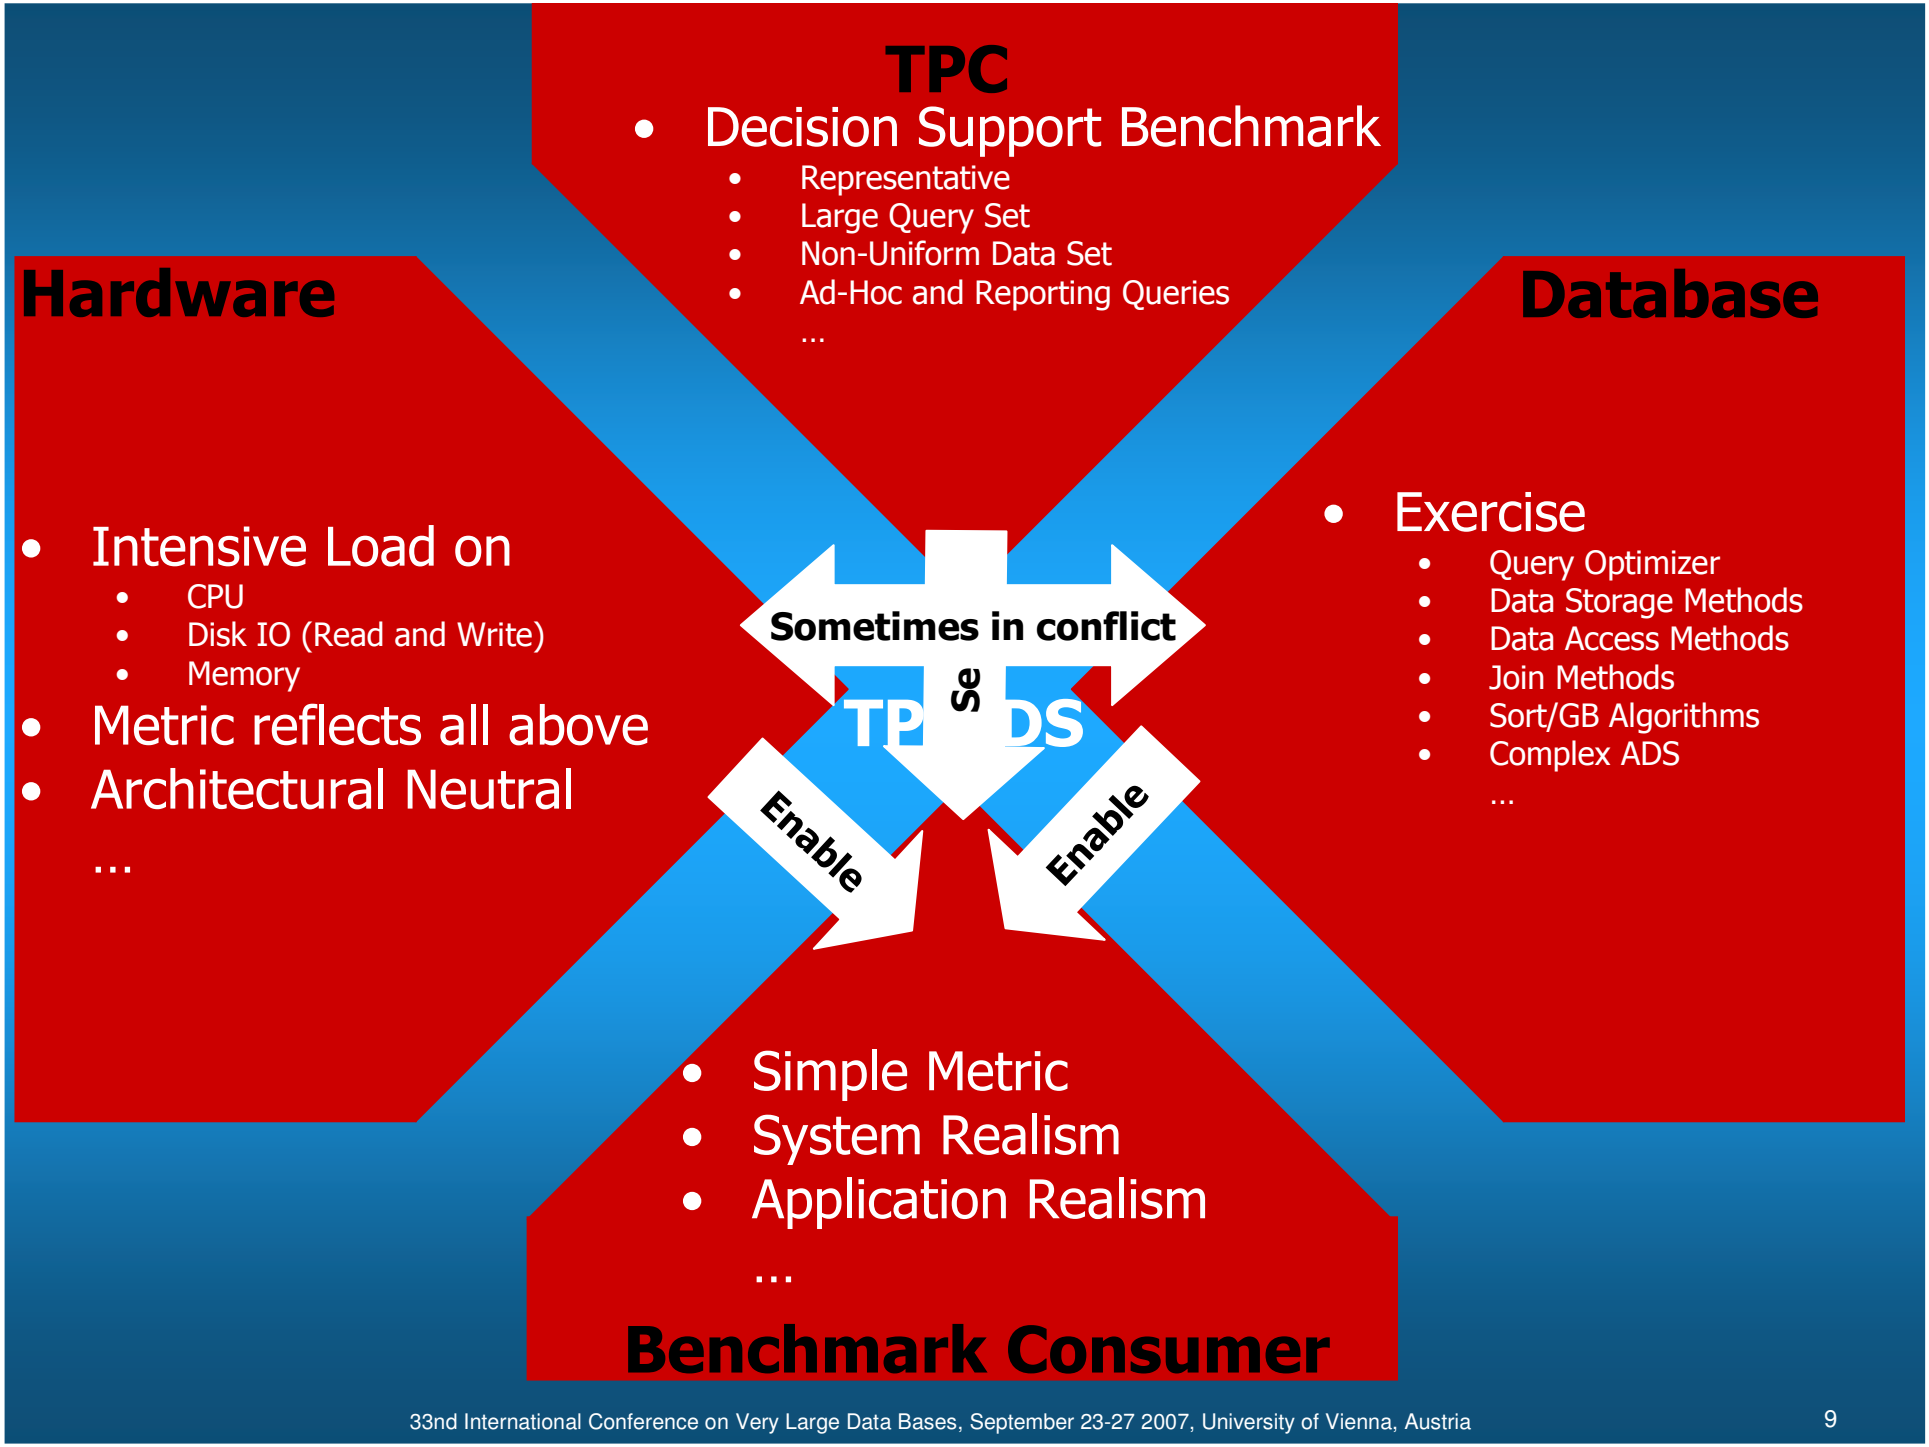
TPC
Decision Support Benchmark
Representative
Large Query Set
Non-Uniform Data Set
Ad-Hoc and Reporting Queries
...
Hardware
Intensive Load on
CPU
Disk IO (Read and Write)
Memory
Metric reflects all above
Architectural Neutral
...
Database
Exercise
Query Optimizer
Data Storage Methods
Data Access Methods
Join Methods
Sort/GB Algorithms
Complex ADS
...
Sometimes in conflict
TPC
DS
Enable
Enable
Simple Metric
System Realism
Application Realism
...
Benchmark Consumer

Benchmark Execution: Bird's Eye View

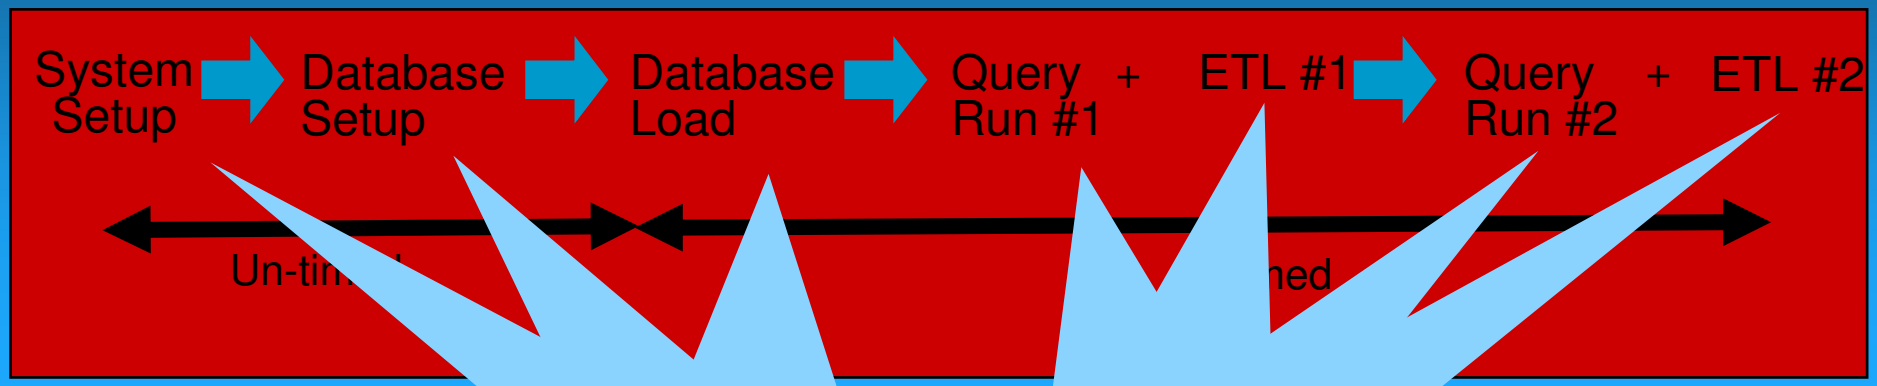

- Data Maintenance-ETL #1
  - Load into fact tables
  - Delete from fact tables
  - Maintain slowly changing dimensions

Hardware Vendor

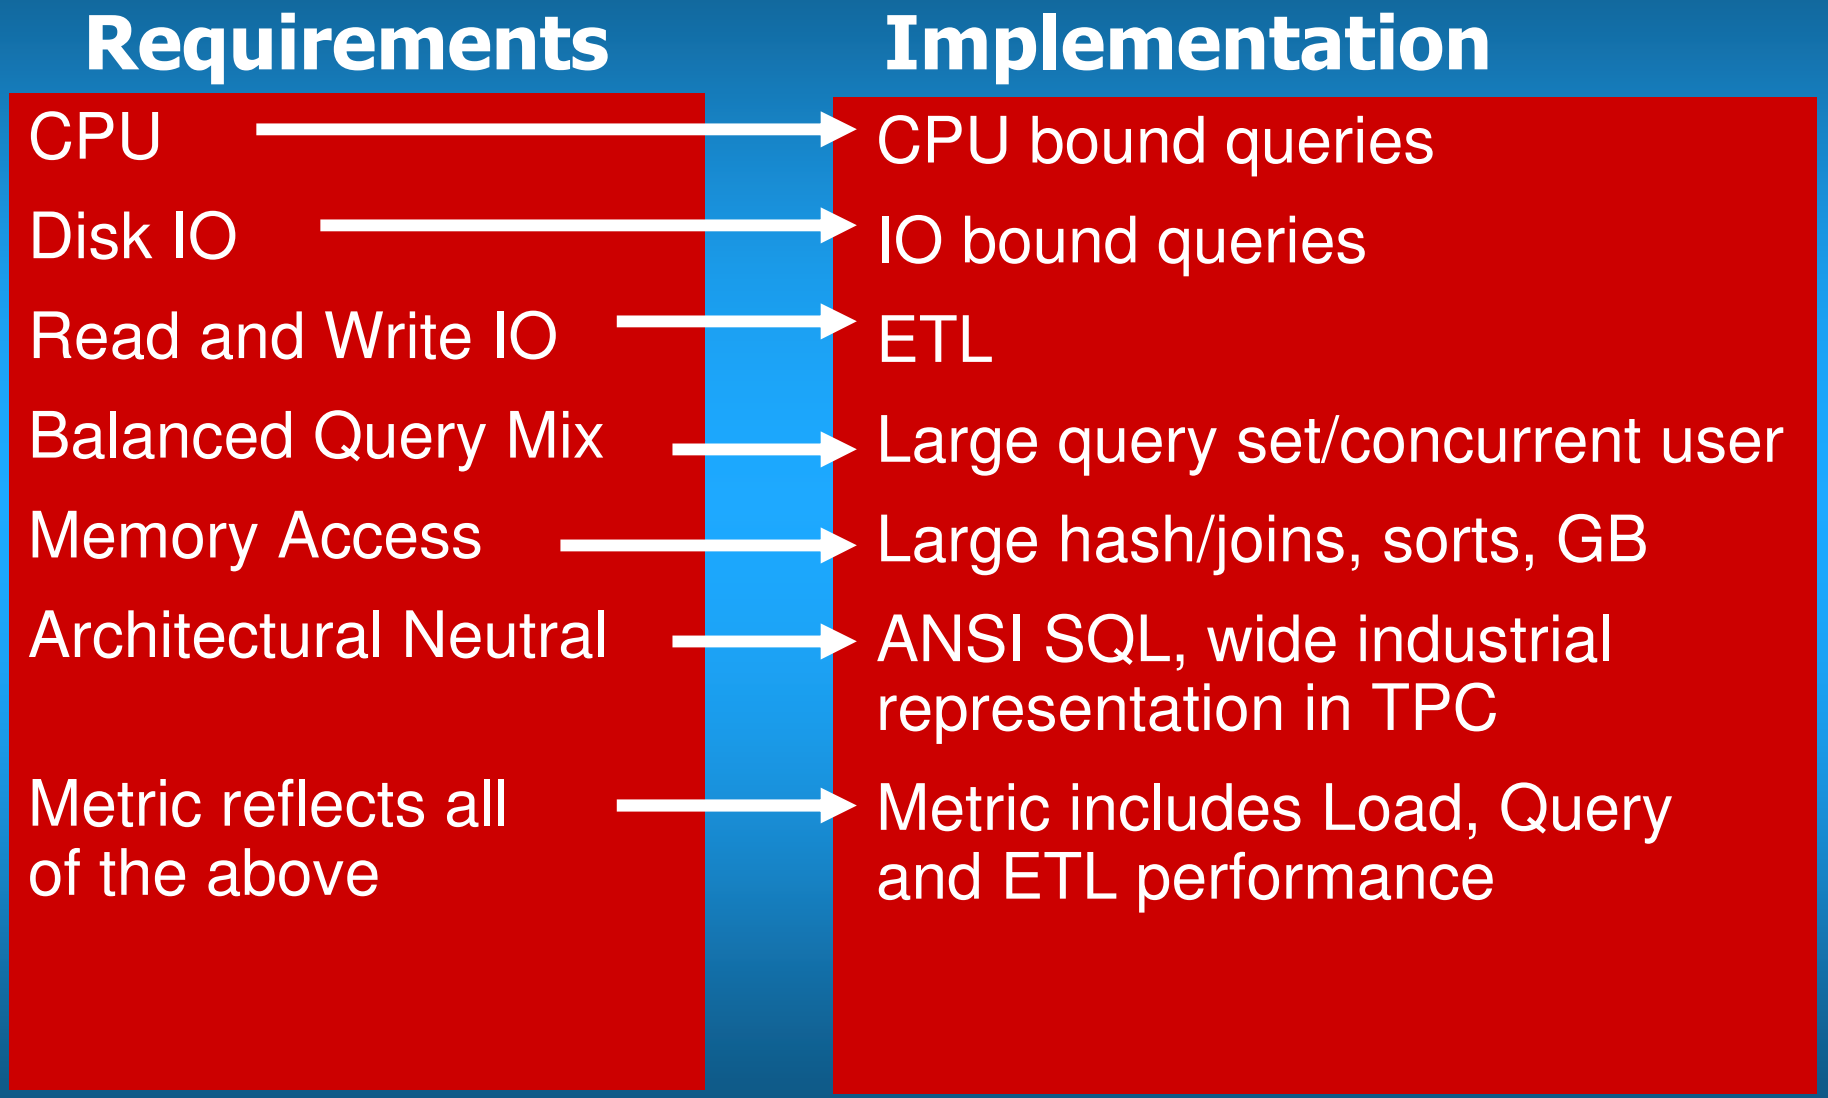

Database Vendor

Requirements	Implementation
Query Optimizer →	Rich Query Set: - star transformation and traditional large join operations Rich Data Set: - NULLs +non-uniform distributions
Join Operations →	Multiple Snowflake Schemas: - Nested Loops - Hash Joins - Bitmap Joins
Sort/GB Operations →	Sort/GB on large data sets
Complex ADS →	ADS are allowed on a subset of the schema
Data Storage Techniques →	Physical Partitioning /Clustering/Compression
Data Access Patterns →	Query Set allows for large sequential scans and random IOs

Query Run

- The query run tests the system's ability to execute the most number of queries in the least amount of time (multi user test)
- Queries can be categorized by:
  - Query Class
    - Ad Hoc
    - Reporting
    - Iterative
    - Data Mining Queries
  - Schema Coverage
  - **Resource Utilization ←**
  - SQL Features

Query Categorization by Resource Utilization

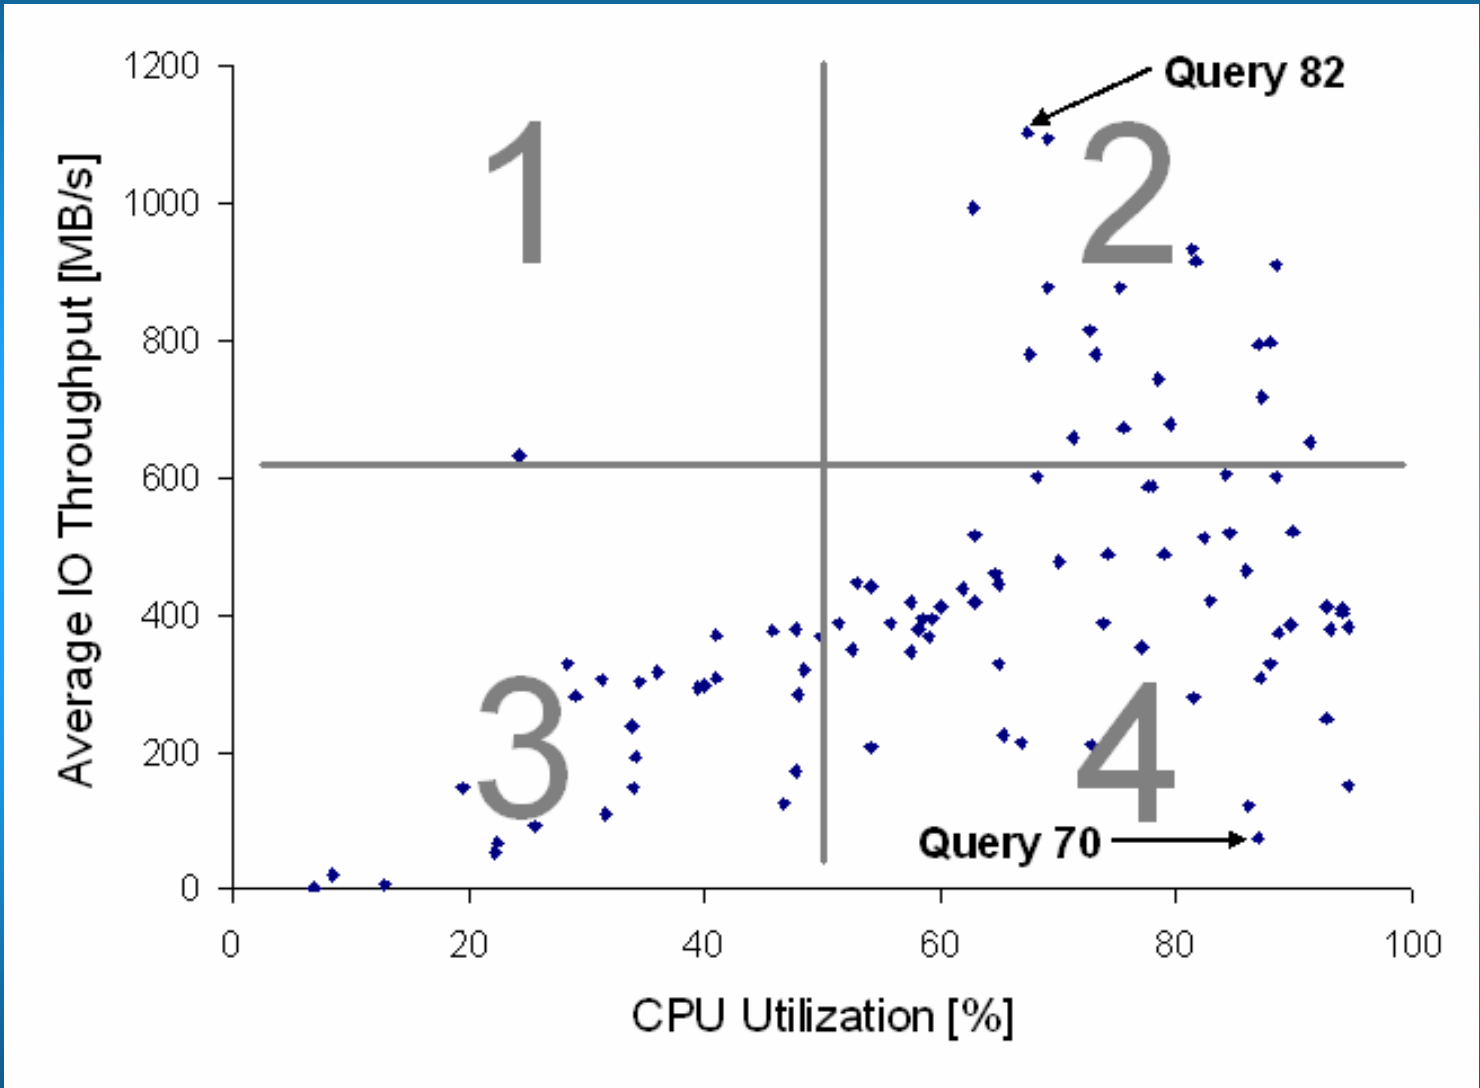

CPU Intensive Query (Query 70)

```
SELECT
  sum(ss_net_profit) as total_sum. s_state,s_county
 ,grouping(s_state)+grouping(s_county)
 ,rank()over(partition by grouping(s_state)
                         +grouping(s_county)
           ,case when grouping(s_county)=0
            then s_state end
  order by sum(ss_net_profit) desc)
FROM store_sales ,date_dim ,store
WHERE d_year = [YEAR]
  AND d_date_sk = ss_sold_date_sk AND s_store_sk  = ss_store_sk
  AND s_state in
     (SELECT s_state
      FROM (SELECT
            s_state ,rank()over(partition by s_state
            order by sum(ss_net_profit)desc) as r
           FROM store_sales,store,date_dim
           WHERE d_year =[YEAR]
             AND d_date_sk = ss_sold_date_sk AND s_store_sk  = ss_store_sk
           GROUP BY s_state)
       WHERE r <= 5)
 GROUP BY ROLLUP(s_state,s_county)
 ORDER BY
   lochierarchy desc
 ,CASE WHEN lochierarchy = 0 THEN s_state END
 ,rank_within_parent;
```

IO Intensive Query (82)

```
SELECT i_item_id
      ,i_item_desc
      ,i_current_price
FROM item, inventory
    ,date_dim ,store_sales
WHERE i_current_price between [P] and [P] + 30
  AND inv_item_sk = i_item_sk
  AND d_date_sk=inv_date_sk
AND d_date between cast('[DATE]' as date)
               AND (cast('[DATE]' as date)+60)
AND i_manufact_id IN ([ID.1],[ID.2],[ID.3])
  AND inv_quantity_on_hand between 100 and 500
  AND ss_item_sk = i_item_sk
GROUP BY i_item_id
        ,i_item_desc
        ,i_current_price
ORDER BY i_item_id;
```

Data Maintenance Functions

- Are defined as pseudo code
- Can be implemented in SQL, or programming SQL
- Need to guarantee referential integrity
- Maintain slowly changing dimensions
- Insert and delete fact tables

Updates/Inserts/Deletes (non history keeping)

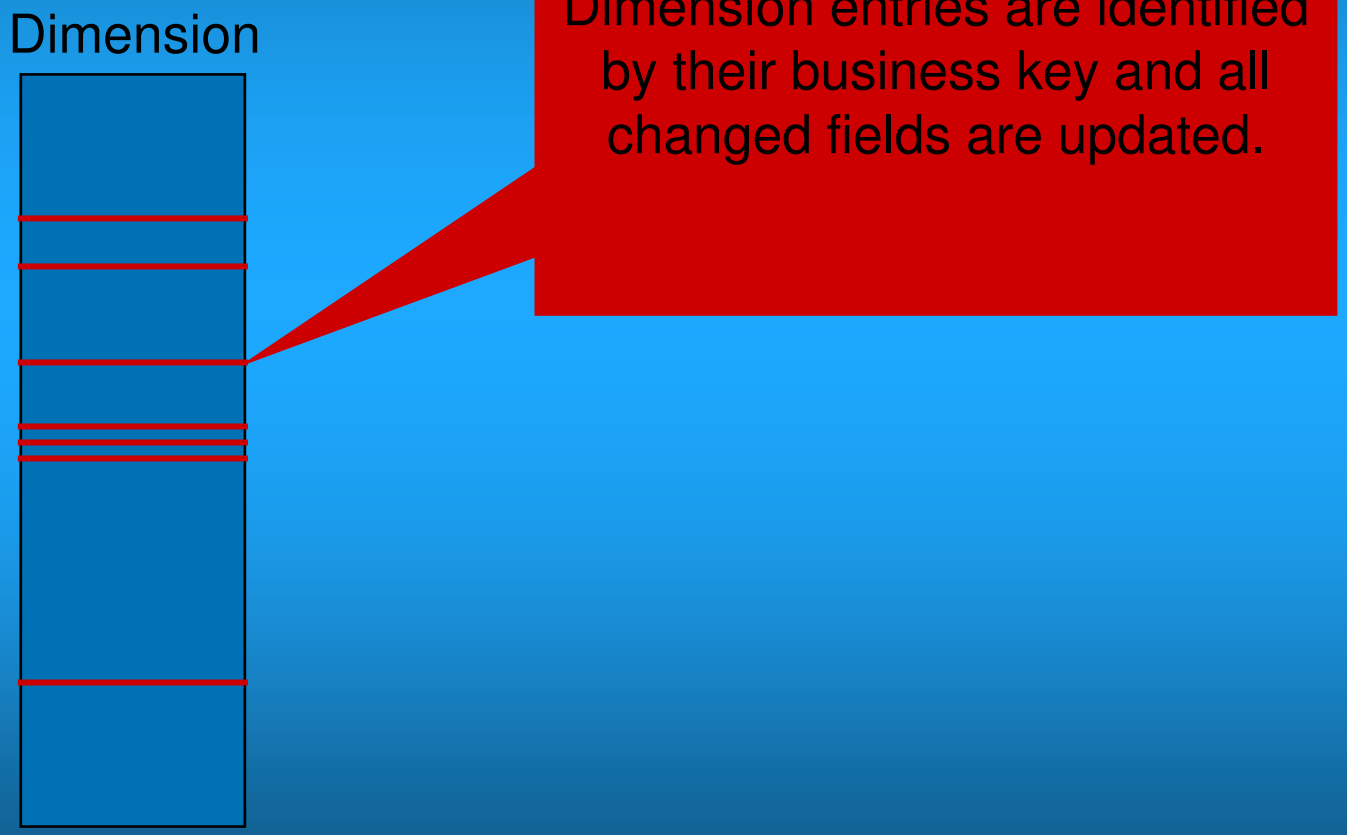

Updates/Inserts/Deletes (history keeping)

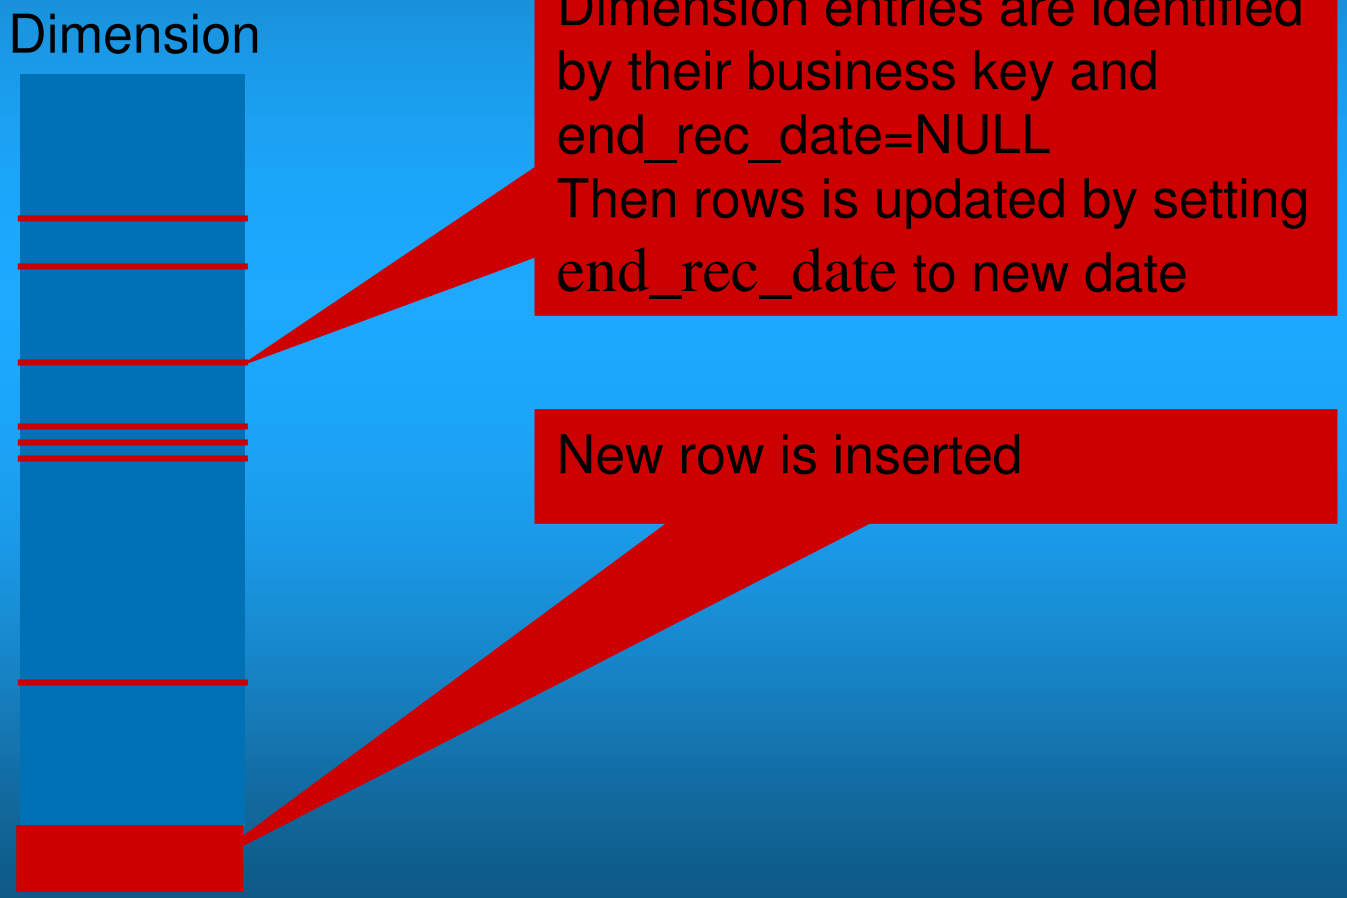

Deletes/Inserts Fact Tables

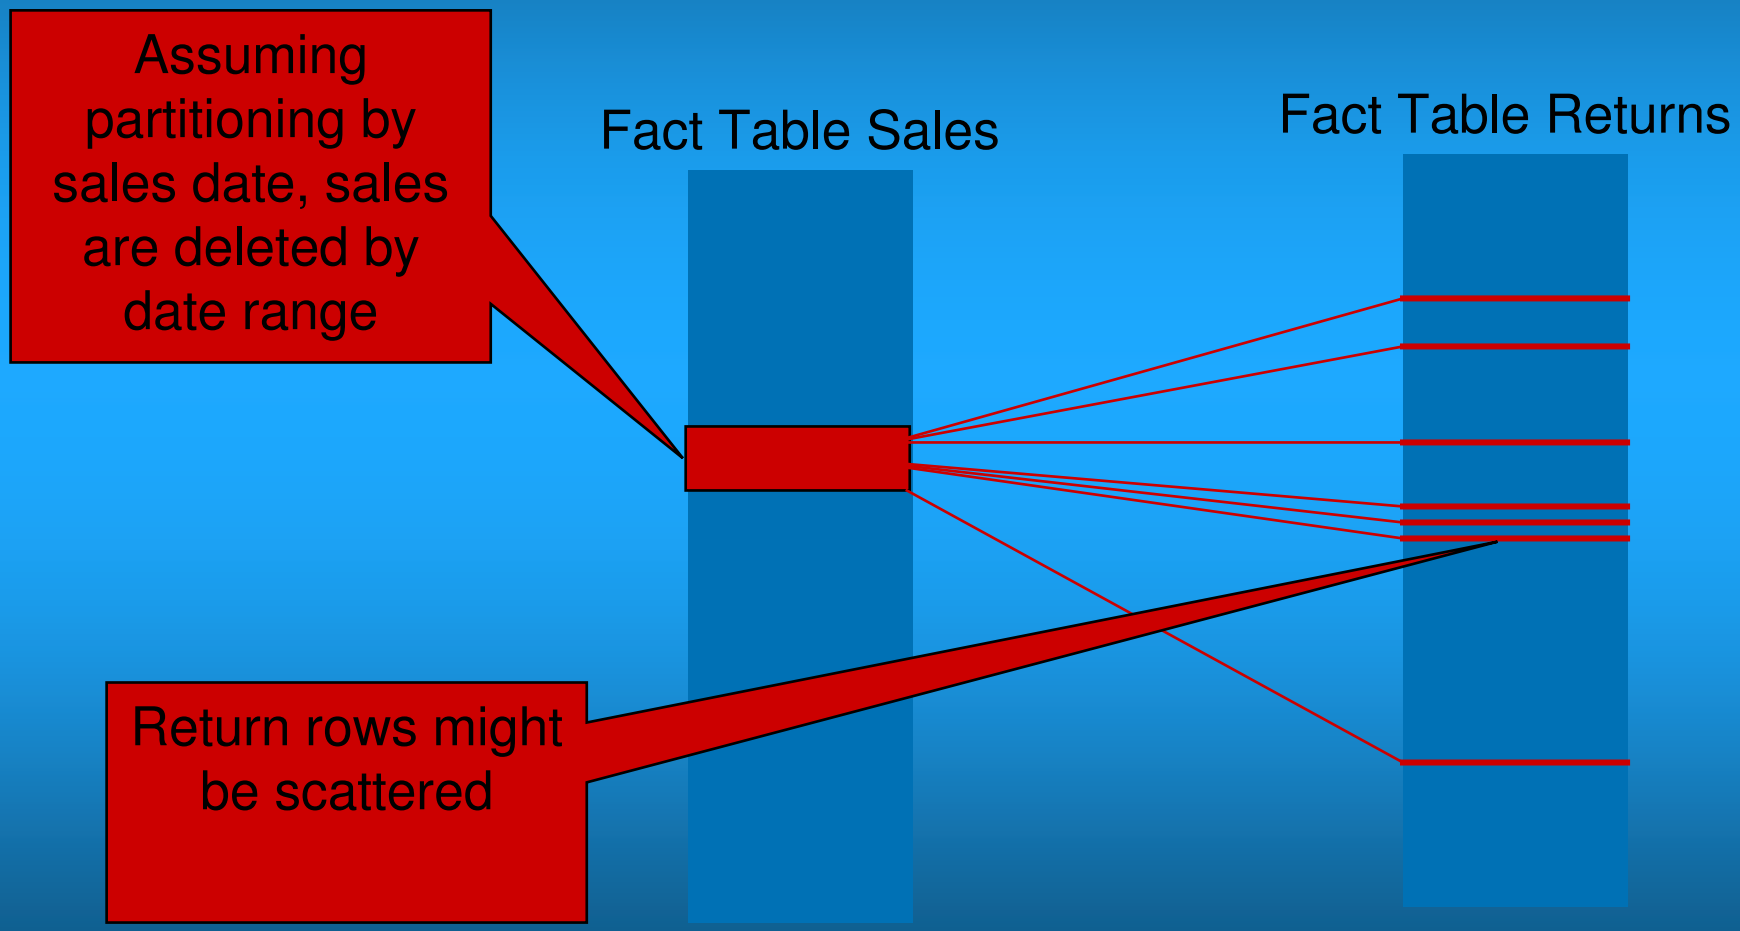

Primary Performance Metric

- Queries per Hour

$$QphDS@SF = \frac{99*2*S*3600*SF}{\left(T_{TT1}+T_{TT2}+0.01*S*T_{Load}\right)}$$

  - S: Number of query streams
  - SF: Scale Factor
  - T_{TT1} and T_{TT2}: elapsed times to complete query run #1 and #2
  - T_{LOAD} is the total elapsed time to complete the database load

Metric Analysis: Use of Materialization

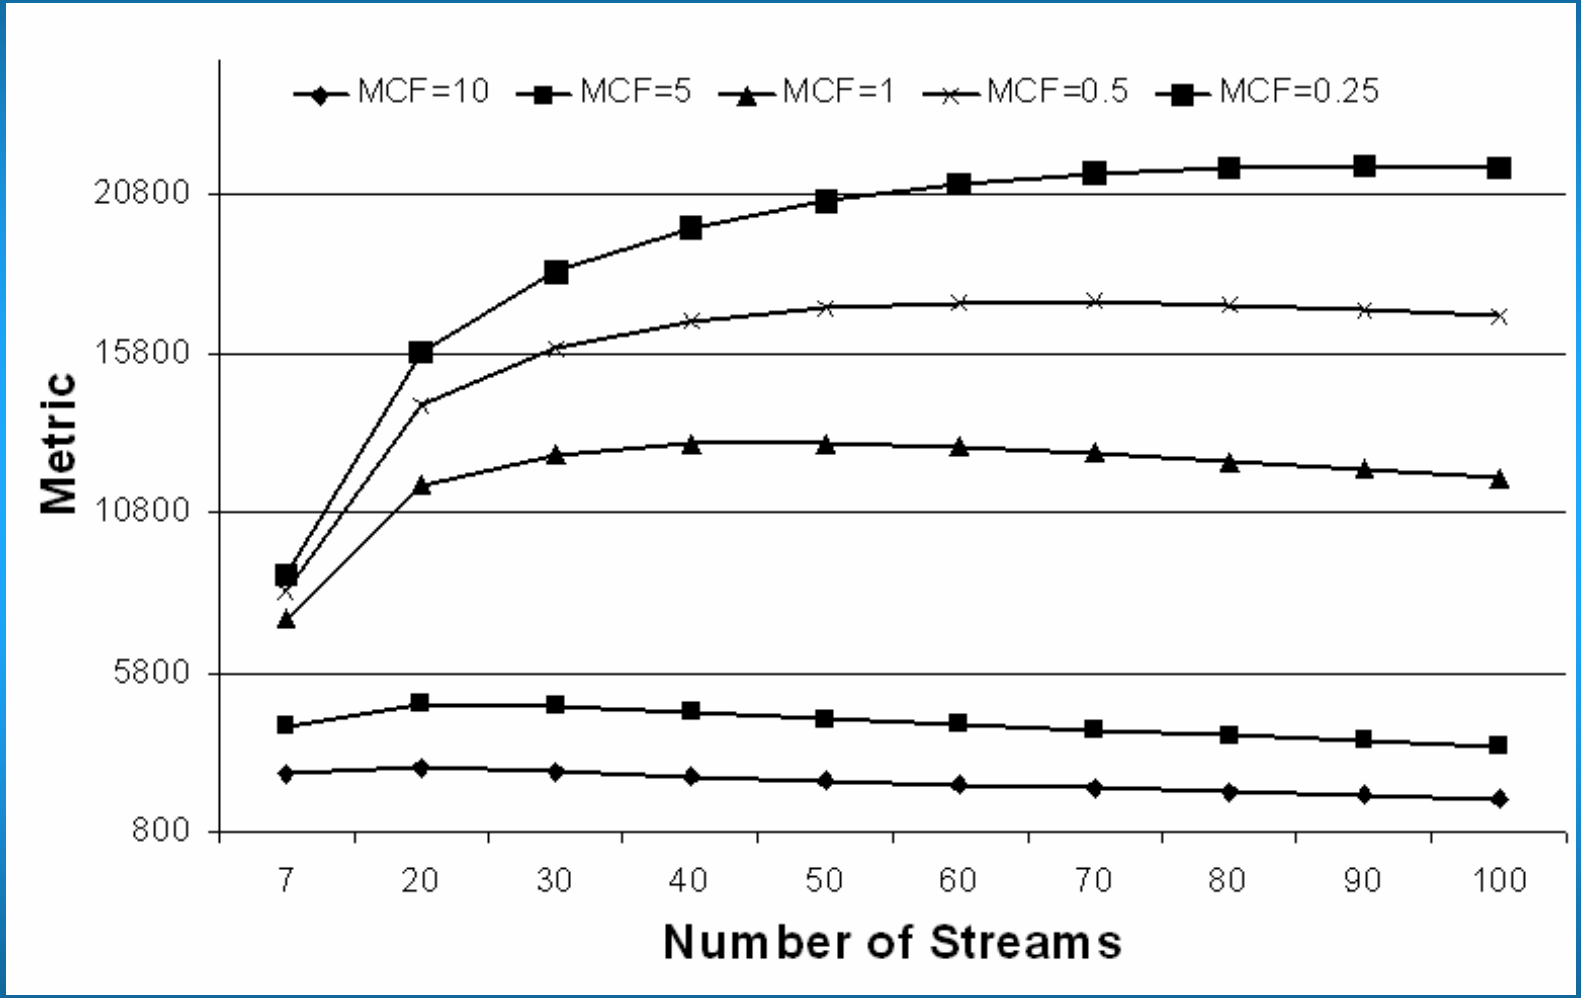

More Information

- Specification
- Dbgen2
- Qgen2
- Query templates

→http://www.tpc.org/tpcds/default.asp

Q & A